Understanding CPU Architecture on Catalyst 4500 E-Series Switches

The Catalyst 4500 and 4900 "E-series" switches are upgraded versions of the "classic series" switches. This document details the CPU packet-handling architecture and outlines some differences between E-series and classic series switches.

Note "E-series" refers to supervisor engines designed for the E+ chassis. These include Supervisor Engine 6-E, Supervisor Engine 6L-E, and Catalyst 4900M. "Classic series" refers to supervisor engines designed for the non E+ chassis. These include WS-X4014 thru WS-X4516-10GE, Catalyst 4948, and Catalyst 4948-10GE.

Note If you are troubleshooting high CPU on Catalyst 4500/4900 switches, please refer to *High CPU Utilization on Cisco IOS Software-Based Catalyst 4500 Switches*.

Note This document applies only to Cisco IOS® Software-based switches and not CatOS-based switches. If you run Catalyst OS (CatOS)-based Catalyst 4500/4000 series switches, refer to the document CPU Utilization on Catalyst 4500/4000, 2948G, 2980G, and 4912G Switches That Run CatOS Software

The information in this document is based on these software and hardware versions:

- Catalyst 4500 E-series switches using supervisor engines Supervisor Engine 6-E or Supervisor Engine 6L-E
- Catalyst 4900 E-series switches including Catalyst 4900M and Catalyst 4948E

This document includes these sections:

Americas Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

CPU Utilization Overview

Hardware-based forwarding switches and Cisco IOS® Software-based routers use the CPU in many different ways. The common misconception is that high CPU utilization indicates a depletion of resources on a device and an imminent crash.

Although one of the symptoms of high CPU utilization is the capacity issue, this issue is almost never a symptom of high CPU utilization with hardware-based forwarding switches like the Catalyst 4500 series switch. The Catalyst 4500 series switch is designed to forward packets in the hardware application-specific integrated circuit (ASIC) and reach traffic-forwarding speeds of up to 250 million IPv4 packets per second (Mpps) on Supervisor Engine 6-E.

The Catalyst 4500 Series Switch CPU performs the following functions:

- Host learning and aging - Host learning and aging in hardware is not supported on Catalyst 4500 series switches and Catalyst 4900 series switches.
- Manages configured software protocols, for example:
  - Spanning Tree Protocol (STP)
  - Routing protocol such as OSPF and EIGRP
  - Hot Standby Routing Protocol (HSRP)
  - Cisco Discovery Protocol (CDP)
  - Port Aggregation Protocol (PAgP)
  - VLAN Trunk Protocol (VTP)
  - Dynamic Trunking Protocol (DTP)
- Programs configuration/dynamic entries to the hardware ASICs, for example:
  - Access control lists (ACLs)
  - CEF entries
- Internally manages various components, for example:
  - Power over Ethernet (PoE) line cards
  - Power supplies
  - Fan tray
- Manages access to the switch, for example:
  - Telnet
  - Console
  - Simple Network Management Protocol (SNMP)
- Forwards packets via the software path, for example:
  - Internetwork Packet Exchange (IPX)-routed packets, which are only supported in the software path
  - Maximum transmission unit (MTU) fragmentation

According to this list, high CPU utilization can result from the receipt or process of packets by the CPU. Some of the packets that are sent for processing might be essential for network operation (for example, bridge protocol data unit (BPDUs) for spanning-tree topology configurations).

Software-forwarded packets require the switching ASICs to send packets to the CPU for processing:

- Packets that are copied to the CPU, but the original packets are switched in hardware, for example:
  - host MAC address learning
- Packets that are sent to the CPU for processing, for example:
  - routing protocol updates
  - BPDUs
  - packets that are sent to the CPU for forwarding (packets that need IPX or AppleTalk routing))

About Catalyst 4500 CPU Packet-Handling Architecture

Similar to the classic series switches, the E-series has an inbuilt quality of service (QoS) mechanism to differentiate between types of traffic that are destined for the CPU. The number of CPU queues is increased to 64 on the E-series switches. Each queue handles various types of packets or events. The following table lists the queues and the packets queued. Because the current software implementation has not yet fully utilized all 64 queues, not all queue numbers are associated with a traffic type. Also, the queue number associated with each of the queue name might be changed slightly between different software releases. Use the **show platform software cpu events** command to check the CPU queue name and number.

Table 0-1 Catalyst 4500 Software Queue Description

Queue Name	Packets Queued
MTU Check Fail	Packets that must be fragmented because the output interface MTU size is smaller than the size of the packet
SaMiss	Frames with unknown source MAC addresses that are copied to the CPU in order to build the Layer 2 forwarding table
PVMapping Miss	Software ensure that this Port VLAN mapping miss does not occur; this queue is not used
Input If Fail	Packets with multicast RPF failure
ESMP	ESMP packets (internal management packets) for the line card ASICs or other component management
L2 Control	Layer 2 control plane packets, such as STP, CDP, PAgP, LACP, or UDLD
Ip Option	Packets with IP header options
Expired Ttl	Routed packets has TTL less than 2
Non Arpa	Packets with an Layer 2 encapsulation other than ARPA
Ucast Rpf Fail	Not used. Default action is to drop.

Table 0-1 Catalyst 4500 Software Queue Description

Queue Name	Packets Queued
SrcIdx Check Fail	Not used.
Adj Same If	Packets that is routed out of the same interface.
RTP	Real Time Protocol packets for voice traffic (Reserved; not used yet)
RSVP	Control packets for voice traffic (Reserved; not used yet)
Input Acl Fwd (Snooping)	Packets that are processed by the DHCP snooping, dynamic ARP inspection, or IGMP snooping features. Control packets that are captured by the input static ACL (such as OSPF, HSRP).
Input Acl Copy (log, unreachable)	Packets that impact an ACE with the log keyword or packets that were dropped because of a "deny" in an input ACL. These packets require the generation of ICMP unreachable messages.
Input Acl Punt	Packets in input direction that are punted to the CPU due to a lack of additional ACL hardware resources, such as TCAM, for security ACL.
Input Acl Err	Not used.
Output Acl Fwd	Not used.
Output Acl Copy	Packets that impact an ACE with the log keyword or packets that in the output direction were dropped due to a deny in an output ACL. These packets require the generation of ICMP unreachable messages.
Output Acl Punt	Packets in output direction that are punted to the CPU because of a lack of additional ACL hardware resources (such as TCAM, for security ACL)
Output Acl Err	Not used.
L2 Bridge	Protocols that are not supported in hardware (such as IPX and AppleTalk routed packets) are bridged to the CPU.
Unknown	Not used.
L2 Router	IPv4 link local range (224.0.0.x) including all Layer 3 protocol such as OSPF, HSRP. IPv6 Link Local Range: (FF02:0:0:0:0:0:XXXX:XXXX - FF02:0:0:0:0:1:XXXX:XXXX)
L3 Glean	Routed packet that need ARP resolution.

Table 0-1 Catalyst 4500 Software Queue Description

Queue Name	Packets Queued
L3 Forward	Packets that must be forwarded in software, such as GRE tunnels.
L3 Receive	Packets with IP destined to router IP address (such as telnet/ssh sessions to router)

Some of the queue names (such as Layer 3 Receive queue) in Table 1 are associated with multiple queue numbers. When packets arrive for a queue type with multiple queues, the packets are placed in the queues based on the QoS label, which is the differentiated services code point (DSCP) value from the IP type of service (ToS).

When a packet that must be processed by the CPU enters a switch, it is first assigned one or multiple CPU events, then placed into the corresponding queue.

Be aware of the following:

- Much of this information is based on Cisco IOS Release 12.2(54)SG. Minor changes might exist between different releases.
- Traffic with different CPU events might be placed into the same queue.
- Traffic with same CPU event might be placed into different queues depending on the priority of the traffic.

In both E-series and classic series switches, the CPU assigns the weights on the basis of importance, or type, and on the basis of traffic priority, or DSCP. The CPU services the queue on the basis of the relative weights of the queue. For example, if both a control packet, such as a BPDU, and an ICMP echo request are pending, the CPU services the control packet first. An excessive amount of low-priority or less-important traffic does not starve the CPU of the ability to process or manage the system. This mechanism guarantees that the network is stable even under high utilization of the CPU. It is essential to understand the ability of the network to remain stable.

If the CPU has already serviced high-priority packets or processes but has spare CPU cycles for some time period, CPU either services the low-priority queue packets or performs background processes of lower priority. High CPU utilization resulting from low-priority packet processing or background processes is considered normal because the software routinely tries to use all the available time. In this way, the CPU strives for maximum performance of the switch and network without compromising switch stability. The Catalyst 4500 series switch considers the CPU underutilized unless the CPU is used at 100 percent for a single time slot.

About CPU Utilization on Catalyst 4500 Classic and 4900 E-Series Switches

Initially, you should identify your switch's CPU utilization with the **show processes cpu** command. The continual update of baseline CPU utilization might be necessary as yoru network setup gets more configuration-rich or as your network traffic pattern changes.

```
Switch# show process cpu sorted
CPU utilization for five seconds: 17%/0%; one minute: 16%; five minutes: 16%
 PID Runtime(ms)   Invoked      uSecs   5Sec   1Min   5Min TTY Process
  60    29111481  13351616       2180  9.27%  8.06%  7.97%   0 Cat4k Mgmt LoPri
  59    25067966  35373845        708  8.00%  8.10%  8.16%   0 Cat4k Mgmt HiPri
  41      323666    430147        752  0.15%  0.14%  0.15%   0 IDB Work
```

```
 131          36       200         180  0.07%  0.01%  0.00%   0 Exec
 107       83017   4301245          19  0.07%  0.05%  0.07%   0 UDLD
 192        5072   3186984           1  0.07%  0.00%  0.00%   0 PM Callback
   6           0         1           0  0.00%  0.00%  0.00%   0 IPC ISSU Receive
   5           0         1           0  0.00%  0.00%  0.00%   0 Retransmission o
   7      340921     58127        5865  0.00%  0.11%  0.06%   0 Check heaps
!--- Output suppressed.
```

This output delineates two processes that use the CPU: *Cat4k Mgmt HiPri* and *Cat4k Mgmt LoPri*). These processes aggregate multiple platform-specific jobs that perform the essential management functions on the switch. These jobs process the control plane and data packets that must be software-switched or processed.

Use the **show platform health** command to observe which of the platform-specific processes use the CPU under the context of *Cat4k Mgmt HiPri* and *Cat4k Mgmt LoPri.*

Each of the platform-specific processes has a target or expected CPU utilization. When that process is within range, the CPU executes the process in the high-priority context and the **show processes cpu** command output displays that utilization under *Cat4k Mgmt HiPri* . If a process exceeds the target, it runs under the low-priority context and the **show processes cpu** command output counts that additional utilization under *Cat4k Mgmt LoPri*. You can also use *Cat4k Mgmt LoPri* to run background and other low-priority processes (such as consistency check and reading interface counters). This mechanism allows the CPU to run high-priority processes when necessary. The remaining idle CPU cycles are used for the low-priority processes. Marginally exceeding the target CPU utilization (or a brief utilization spike) does not reflect a problem that requires investigation.

```
Switch# show platform health | exc 0.00
                     %CPU   %CPU    RunTimeMax   Priority  Average %CPU  Total
                     Target Actual Target Actual   Fg   Bg 5Sec Min Hour  CPU
VSI channel Slot-01    1.00   0.27      6      1  100  500    0   0    0  17:45
VSI channel Slot-04    1.00   0.05      6      0  100  500    0   0    0  5:40
VSI channel Local Ja   1.00   0.01      6      0  100  500    0   0    0  3:28
VSI channel Remote J   1.00   0.04      6      0  100  500    0   0    0  3:16
GalChassisVp-review    3.00   3.79     10    157  100  500    3   3    3  233:53
Lj-poll                1.00   0.01      2      0  100  500    0   0    0  1:29
StatValueMan Update    1.00   0.07      1      0  100  500    0   0    0  4:05
GalK5TatooineStatsMa   0.70   0.02      4      0  100  500    0   0    0  2:12
K5L3FlcMan Consisten   2.00   0.48     15      7  100  500    0   0    0  55:28
K5L3FlcMan NI Regs &   1.00   0.42      5      4  100  500    0   0    0  22:51
K5L3AdjStatsMan Revi   2.00   0.03     10      6  100  500    0   0    0  15:00
K5FlcHitMan review     2.00   0.01      5      2  100  500    0   0    0  20:41
K5PortMan Regular Re   2.00   0.15     15     11  100  500    0   0    0  21:35
K5PortMan Ondemand L   3.00   0.34     30      0  100  500    0   0    0  18:51
K5 L2 Aging Table Re   2.00   0.04     20      4  100  500    0   0    0  8:55
K5ForerunnerPacketMa   1.50   0.39      4      0  100  500    0   0    0  37:00
K5ForerunnerPacketMa   0.70   0.24      4      0  100  500    0   0    0  15:56
K5QosDhmMan Rate DBL   2.00   4.26      7      5  100  500    4   4    4  327:28
K5VlanStatsReview      2.00   0.82     10      4  100  500    0   0    0  67:35
K5AclCamMan Audit re   1.00   0.06     10      5  100  500    0   0    0  12:49
K5AclCamStatsMan hw    3.00   0.14     10      5  100  500    0   0    0  13:32
RkiosPortMan Port Re   2.00   0.15     12      4  100  500    0   0    0  17:14
Rkios Module State R   4.00   0.02     40      2  100  500    0   0    0  1:39
Rkios Online Diag Re   4.00   0.01     40      0  100  500    0   0    0  1:34
RkiosIpPbr IrmPort R   2.00   0.02     10      1  100  500    0   0    0  1:41
RkiosAclMan Review     3.00   0.03     30      0  100  500    0   0    0  3:02
Xgstub Stats Review    0.50   0.09      5      0  100  500    0   0    0  6:04
EthHoleLinecardMan(2   1.11   1.52     10      3  100  500    1   2    1  137:30
Xgstub Stats Review    0.50   0.10      5      0  100  500    0   0    0  7:08
```

```
EthPoeControllerMan(   0.20   0.01      2      0  100  500    0   0    0  0:06
EthPoeControllerMan(   0.20   0.01      2      0  100  500    0   0    0  0:07
Xgstub Stats Review    0.50   0.09      5      0  100  500    0   0    0  6:09
Xgstub Stats Review    0.50   0.09      5      0  100  500    0   0    0  6:03
EpmPortGroup(0:N) on   0.50   0.01      4      0  100  500    0   0    0  0:24
EthHoleLinecardMan(8   1.11   2.38     10      3  100  500    1   2    1  75:56
                     -------------
%CPU Totals          238.72  22.42
```

About the show platform health Command

The **show platform health** command provides information that is particularly relevant to a development engineer. To troubleshoot high CPU utilization, look for a higher number in the **%CPU actual** column in the command's output. To verify the CPU usage of that process, observe the **1 minute** and **1 hour average %CPU** columns on the right side of the row. Sometimes, processes momentarily peak but do not hold the CPU for a long time. Some of the momentarily high CPU utilization occurs during hardware programming or programming optimization. For example, a spike of CPU utilization is typical during the hardware programming of a large ACL in the TCAM.

The following table provides some basic information about the important platform-specific processes that appear in the output of the **show platform health** command. If you run into high CPU issue with any platform dependent process other than *K5CpuMan Review*, contact Cisco Technical Assistance Center (TAC) for support

Table 0-2 Catalyst 4500 Platform-Specific Process Names

Platform-Specific Process Name	Description
VSI Channel	Linecard-to-Supervisor/Supervisor-to-supervisor communication process
GalChassisVp-review Pim-review	Monitoring various state of the chassis and also includes some line cards states and PoE
S2w-JobEventSchedule	Manages the S2W protocols to monitor line cards state
Stub-JobEventSchedul	Stub ASIC-based line card monitoring and maintenance
Pim-review	Chassis or line card state management
Ebm *process name*	Ethernet bridge module, such as aging and monitoring
KxAclPathMan	ACL state management and maintenance
K5L3 *process name*	Different Layer 3 process that managing various Layer 3 function and tables such as forwarding entries, adjacencies, multicast RET entries, and statistics
K5PortMan *process name*	Manages the various port-related programming functions such as status review, update, statistics, and tx-queues
K5L2 *process name*	Various Layer 2 processes are responsible for maintenance of the various Layer 2 tables such as MAC address table, aging time, and Layer 2 multicast table.
K5RetStatsMan Review	Manages (Replication Expantion Table) RET statistics

Table 0-2 Catalyst 4500 Platform-Specific Process Names

Platform-Specific Process Name	Description
K5CpuMan Review	The process that performs software packet forwarding. This job also en-queues and extracts packets from CPU packet queues. If you see high CPU utilization due to this process, this typically indicates that the high CPU is caused by traffic.
K5QosPolicerStatsMan	Manages the QoS policer statistics
K5QosDhmMan Rate DBL	Manages QoS DBL
K5VlanStats *process name*	Manages and review VLAN statistics
K5Acl *process name*	Manages updates or reviews the input and output TCAM hardware for QoS and security ACL and snooping features
RkiosPortMan Port Review	Port state monitoring and maintenance
Rkios Module State Review	Line card monitoring and maintenance
EthHoleLinecardMa	Manages GBICs in each of the line cards
Quack	The process that checks line card authenticity. One process will be generated for each linecard.
EthPhyControllerMan	Manages and control the PHY on E-series line cards
EthPoeControllerMan	Manages and control the PoE on E-series line cards
Xgstub Stats Review	Review the statistic on the E-series line card stub ASICs

Control Traffic Interception Mechanism

On a Catalyst 4500 series switch, all Layer 2 and Layer 3 control packets are captured in the CPU by a static access-list, which is automatically programmed during boot up. Use the **show platform hardware acl input entry static** command to check this access-list as well as the number of hits. For some of the protocols (OSPF and HSRP), the access-list cannot determine whether the feature is configured; all the control traffic is captured. For example, in a U or V shape topology, the HSRP hello packets between the two upstream routers are captured to the CPU on the Catalyst 4500 switch even if HSRP is not configured. This implementation might create high CPU problems in some corner cases. (See *High CPU Utilization on Cisco IOS Software-Based Catalyst 4500 Switches* for more details.) Because classic and non-classic series supervisors use the same implementation, all the supervisor engines experience the same problem.

```
Switch#  show platform hardware acl input entries stat

    Input Static ACL Cam Entries

BlockId: 30, LookupType: Security, BlockWidth: 320Bit

     CamIndex Entry Type                        Active       Hit Count
     -------- ----------                        ------       ---------
        61440 DenyIpv6SrcAddrLoopback             Y                0 (estimate)
        61442 DenyIpv6SrcAddrLoopback             Y                0 (estimate)
        61444 DenyIpv6SrcAddrMcast                Y                0 (estimate)
        61446 DenyIpv6SrcAddrMcast                Y                0 (estimate)
!--- Output suppressed.
        61498 PermitIpv6LinkLocalNdAd             Y                0 (estimate)
        61500 Ipv6PuntToCpu                       Y                0 (estimate)
        61502 Ipv6PermitAll                       Y                0 (estimate)
```

```
        63486 Ipv6EndOfCam                            Y                0 (estimate)

BlockId: 31, LookupType: Security, BlockWidth: 160Bit

     CamIndex Entry Type                           Active       Hit Count
     -------- ----------                           ------       ---------
        63488 CaptureInputEsmp                       Y              380 (estimate)
        63489 CaptureEapol                           Y                0 (estimate)
        63490 DropDot1dFlowControl                   Y                0 (estimate)
        63491 PermitLldp                             Y                0 (estimate)
        63492 PermitBpdusRange1                      Y              377 (estimate)
        63493 DropBpdus                              Y                0 (estimate)
        63494 DropBpdusRange2(02-03)                 Y                0 (estimate)
        63495 DropBpdusRange3(04-07)                 Y                0 (estimate)
        63496 DropBpdusRange4(08-0F)                 Y                0 (estimate)
        63497 PermitCdp                              Y              182 (estimate)
        63498 CopyHfl                                Y                0 (estimate)
        63499 CapturePppoeDiscovery                  N                0 (estimate)
        63500 CopyMmu                                Y                0 (estimate)
        63501 PermitGarp                             Y                0 (estimate)
        63502 DropGarp                               Y                0 (estimate)
        63503 PermitSharedStp                        Y                1 (estimate)
        63504 PermitLoopbackTest                     Y                2 (estimate)
        63505 DenyIsl                                Y                0 (estimate)
        63506 DenyIsl                                Y                0 (estimate)
        63507 PermitProtTunnel                       N                0 (estimate)
        63508 DenyMulticastSource                    Y                0 (estimate)
        63509 CaptureCgmp                            Y                0 (estimate)
        63510 CaptureOspf                            Y                0 (estimate)
        63511 CaptureIgmp                            Y                0 (estimate)
        63512 CapturePim                             Y                0 (estimate)
        63513 CaptureHsrpV2                          Y                0 (estimate)
        63514 CaptureAllSystemsOnSubnet              Y                0 (estimate)
        63515 CaptureAllRoutersOnSubnet              Y               41 (estimate)
        63516 CaptureRipV2                           Y                0 (estimate)
        63517 CaptureRsvdMcastAddressRange           Y                0 (estimate)
        63518 CaptureDhcpClientToServer              N                0 (estimate)
        63519 CaptureDhcpServerToClient              N                0 (estimate)
        63520 CaptureDhcpServerToServer              N                0 (estimate)
        63521 Ipv4HeaderException                    Y                0 (estimate)
        63522 Ipv6HeaderException                    Y                0 (estimate)
        63523 MartianIp                              Y                0 (estimate)
        63524 PuntToCpu                              Y                0 (estimate)
        63525 PermitAll                              Y              273 (estimate)
        65535 EndOfCam                               Y                0 (estimate)
```

Troubleshooting Common High CPU Utilization Problems

This section covers some of the common high CPU utilization problems on the Catalyst 4500 switches

Topics include:

High CPU Utilization Due to Process-Switched Packets

Please refer to *High CPU Utilization on Cisco IOS Software-Based Catalyst 4500 Switches* for details. Except for the differences explained earlier in this document, the troubleshooting steps are very similar.

High CPU Utilization Due to K5L3 Review Jobs with Incomplete ARPs

An E-series Catalyst 4500 switch can display high CPU utilization in several K5L3 review processes (K5L3FlcMan FwdEntry, K5L3Unciast IFE Review, and K5L3UnicastRpf IFE Review). When a switch receives a packet with an unknown local destination IP address, it sends an ARP request to the attached subnet. Simultaneously, the switch reviews and updates the routing table. High CPU could be triggered either by a small amount of arp (10 packets) with a large routing table (more than 100K+ routes) or a large amount of arp (thousands of packets/sec) with a small routing table (1K routes). This behavior was optimized in Cisco IOS software Releases 12.2(50)SG6, 12.2(53)SG1, or later. (CSCta77487)

```
Switch# show platform health
                    %CPU   %CPU    RunTimeMax   Priority  Average %CPU  Total
                    Target Actual Target Actual   Fg   Bg 5Sec Min Hour  CPU
!--- Output suppressed.
K5L3FlcMan FwdEntry    2.00  27.51     15     14  100  500   25  26   20  4454:02
K5L3Unciast IFE Revi   2.00  31.28     15     10  100  500   26  26   21  4695:14
K5L3UnicastRpf IFE R   2.00  31.41     15      7  100  500   26  26   20  4659:17
```

High CPU Utilization Due to RSPAN

An E-series Catalyst 4500 switch configured with RSPAN might display high CPU utilization during the host learning process. Although the RSPAN VLAN does not learn any MAC addresses, a copy of the packet with unknown source MAC is still sent to and dropped at the CPU.

```
Switch# show processes cpu sorted
CPU utilization for five seconds: 93%/7%; one minute: 94%; five minutes: 96%
 PID Runtime(ms)   Invoked      uSecs   5Sec   1Min   5Min TTY Process
  49   2095141161223088784        171 84.39% 84.85% 87.07%   0 Cat4k Mgmt LoPri
  48     1195120   4781112        249  1.91%  1.86%  1.84%   0 Cat4k Mgmt HiPri
!--- Output suppressed

Switch# show platform cpu packet statistics all
!--- Output suppressed
Packets Dropped In Processing by CPU event

Event             Total                5 sec avg 1 min avg 5 min avg 1 hour avg
----------------- -------------------- --------- --------- --------- ----------
Unknown                             0         0         0         0          0
Sa Miss                    2600617361     17399     15937     12797      12257
```

The behavior changed in Cisco IOS Releases 12.2(50)SG4, 12.2(53)SG, or later; packets are processed in hardware and no extra copy is sent to the CPU. (CSCsu81046)

High CPU Due to K5AclCamStatsMan

The E-series Catalyst 4500 switches can display high CPU utilization in the K5AclCamStatsMan hw process whena large number of ACL entries are configured with hardware statistics enabled. By default, certain applications (IP Source Guard and QoS) enable ACL statistics, triggering high CPU.

```
Switch# show processes cpu sorted
CPU utilization for five seconds: 75%/0%; one minute: 80%; five minutes: 79%
 PID Runtime(ms)   Invoked      uSecs   5Sec   1Min   5Min TTY Process
  54    16177852   1680243       9628 58.89% 62.61% 63.19%   0 Cat4k Mgmt LoPri
  53     3156396   2829272       1115 11.65% 11.49% 11.39%   0 Cat4k Mgmt HiPri

Switch# show platform health
                     %CPU   %CPU    RunTimeMax   Priority  Average %CPU  Total
                     Target Actual Target Actual   Fg   Bg 5Sec Min Hour  CPU
!--- Output suppressed.
K5AclCamMan Audit re   1.00   3.00     10      6  100  500    2   2    1  10:38
K5AclCamStatsMan hw    3.00  38.26     10     14  100  500   51  56   44  177:53
```

This behavior was optimized in Cisco IOS Releases 12.2(50)SG6, 12.2(53)SG1, or later. If the QoS configuration caused the high CPU utilization, you need both a software update and the **no qos statistics classification** command to reduce the consumption. (CSCta54369)

High CPU Utilization from Control Traffic Interception

Capturing Layer 2 and Layer 3 control packets using a static ACL might cause high CPU problems in some U or V shape topologies. Consider the network topology below:

ILLO

Two distribution layer switches are configured as Layer 2 or Layer 3 and the access layer consists of Catalyst 4500 switches configured as pure Layer 2 devices. The distribution switches run Layer 3 protocols (like HSRP and OSPF) and send keepalives. This Layer 3 control traffic would pass through the access layer Catalyst 4500 series switch. Even though an access switch is disinterested in the Layer 3 control traffic, the traffic is still punted to CPU because of th control interception mechanism. The CPU will perform the look up and subsequently forward the traffic. This will increase the baseline CPU usage but not pose a problem.

These are factors that determine the CPU baseline usage from passing Layer 3 control traffic:

- The number of Layer 3 protocols configured
- The number of VLANs running the Layer 3 protocols
- The length of keepalive or hello interval
- The number of interfaces in the VLAN that that protocol packets are being forwarded out to.

When the amount of Layer 3 control traffic is large and causes high CPU, the Layer 3 protocol might flap due to delay or to dropping the keepalive packets.

The default behavior can be changed by configuring the control packet capture mode.

```
Switch(config)# access-list hardware capture mode vlan
```

Refer to the "Selecting Mode of Capturing Control Packets" in the Catalyst 4500 Series Switch Software Configuration Guide for guideline and restrictions.

Troubleshooting Tools to Analyze the Traffic Destined to the CPU

All E-series switches have CPU analyzing tools such as the CPU traffic SPAN and the Built-In CPU sniffer. Refer to the High CPU troubleshooting document from corresponding information for the classic supervisor engines.

Summary

The Catalyst 4500 E-Series switches have similar but improved CPU architecture relative to non-classic switches. E-series switches can handle a high rate of both IPv4 and IPv6 packet forwarding in hardware. Some of the features or exceptions can cause the forwarding of some packets through the CPU process path.

The Catalyst 4500 E-series supervisor engine has an improved QoS mechanism to handle CPU-bound packets relative to non-classic switches. This mechanism ensures reliability and stability of the switches and maximizes the CPU for packet forwarding in software.

Similar to the classic switches, the E-series switches, has the powerful tools and sufficient commands to aid in the identification of the root cause of high CPU-utilization scenarios.

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html

Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS Version 2.0.